

IPEX

Traffic Measurements

CNRI:

C. Brownstein

Telcordia Technologies:

K. R. Krishnan, Marc Pucci,
Allen McIntosh, and Chungmin Chen

SLAC:

Connie Logg, Les Cottrell, and
Warren Matthews

DARPA NMS Meeting
Baltimore, April 18, 2002

IPEX Phase I

- **Goals (Telcordia's contract with CNRI)**
 - Collect samples of source traffic of commercial sites continually for use by NMS community
 - Establish low maintenance / high-availability network of measurement servers for monitoring and real-time experimentation

Deployment

- **Data Collectors**

- Operational at Telcordia, SLAC, and West Group Publishing; Kaiser Permanente next on list
- Equipment for two additional sites configured and available for installation

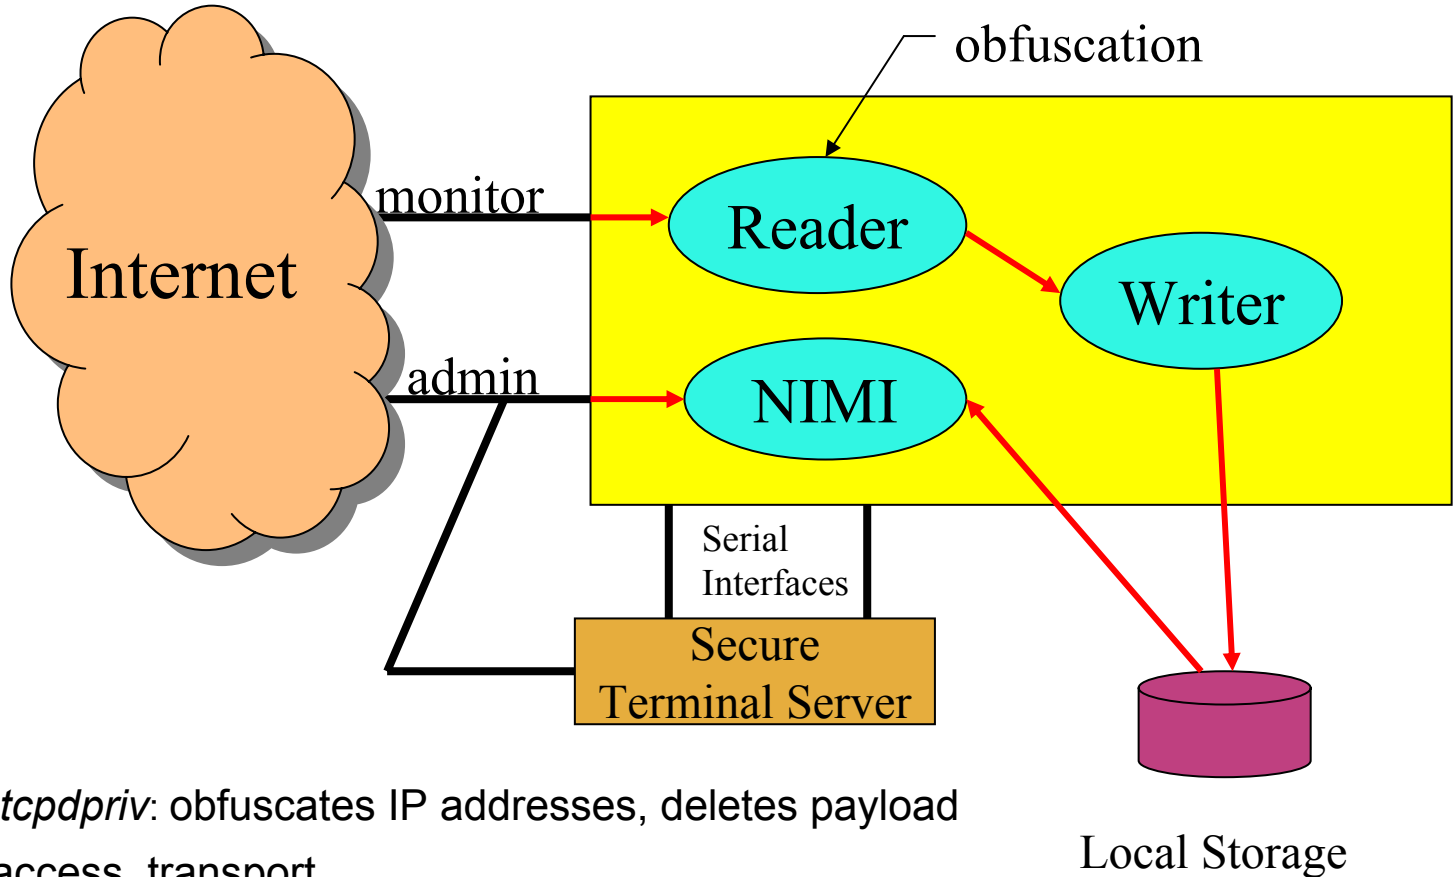
- **Traffic Data**

- Using limited CPU cycles, simple statistical summaries (packet sizes, applications) gathered and continuously shipped back to central Telcordia site
- Detailed trace for busy half-hour also shipped back
- Other traces available
- Statistics stored in database

IPEX Infrastructure Capabilities

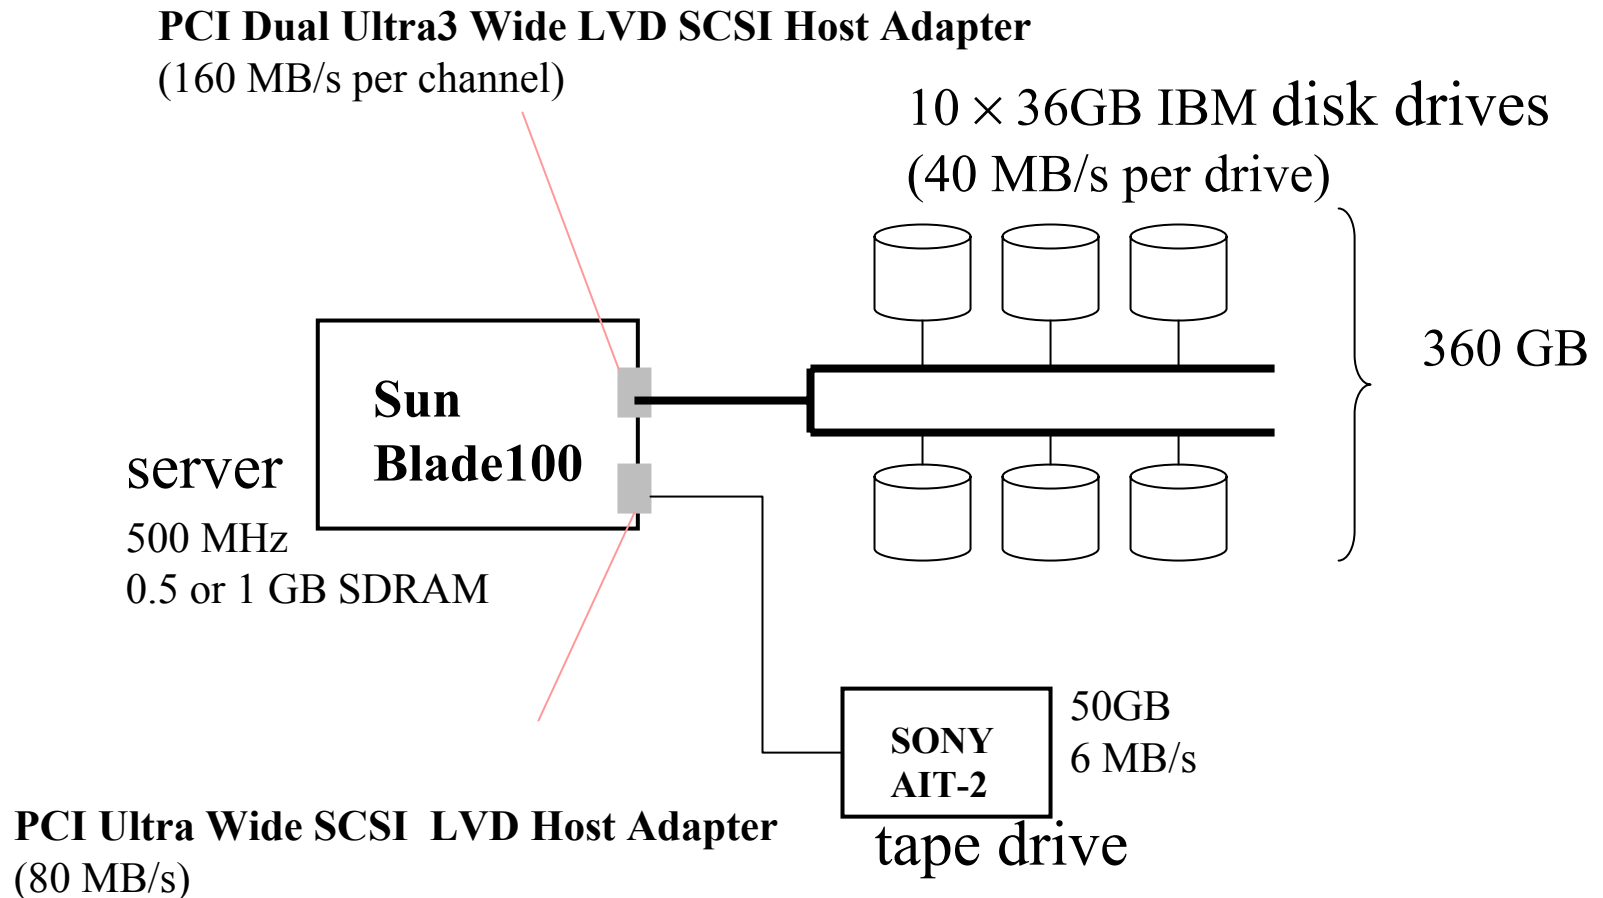
- Central management service (from Telcordia)
 - Achieves “lights out” operation (i.e., remote start and re-start) of all sites
 - Schedules and co-ordinates collection and reporting of data from individual sites (*particularly valuable for conducting controlled experiments in next phase of IPEX*)
 - *Coming*: web interface for access to data traces and summaries

Data Collector Architecture

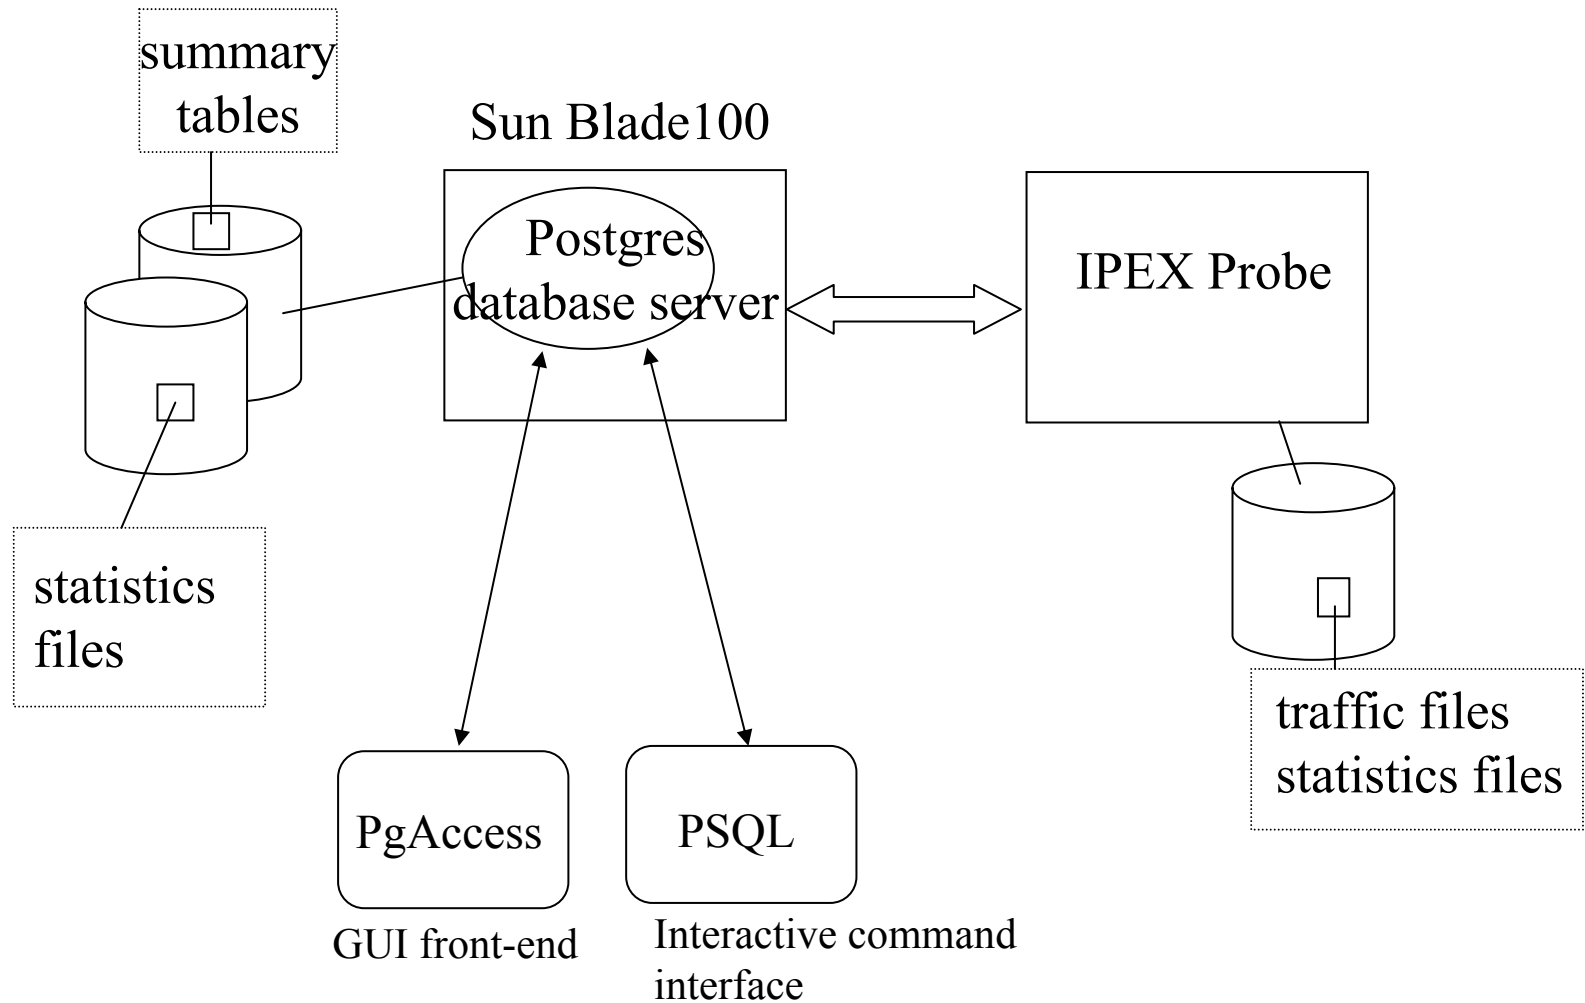


- Modified *tcpdpriv*: obfuscates IP addresses, deletes payload
- NIMI for access, transport
- Deployed on Sparc Netra T1-100, 200
- Runs under Linux, FreeBSD

Data Storage Architecture



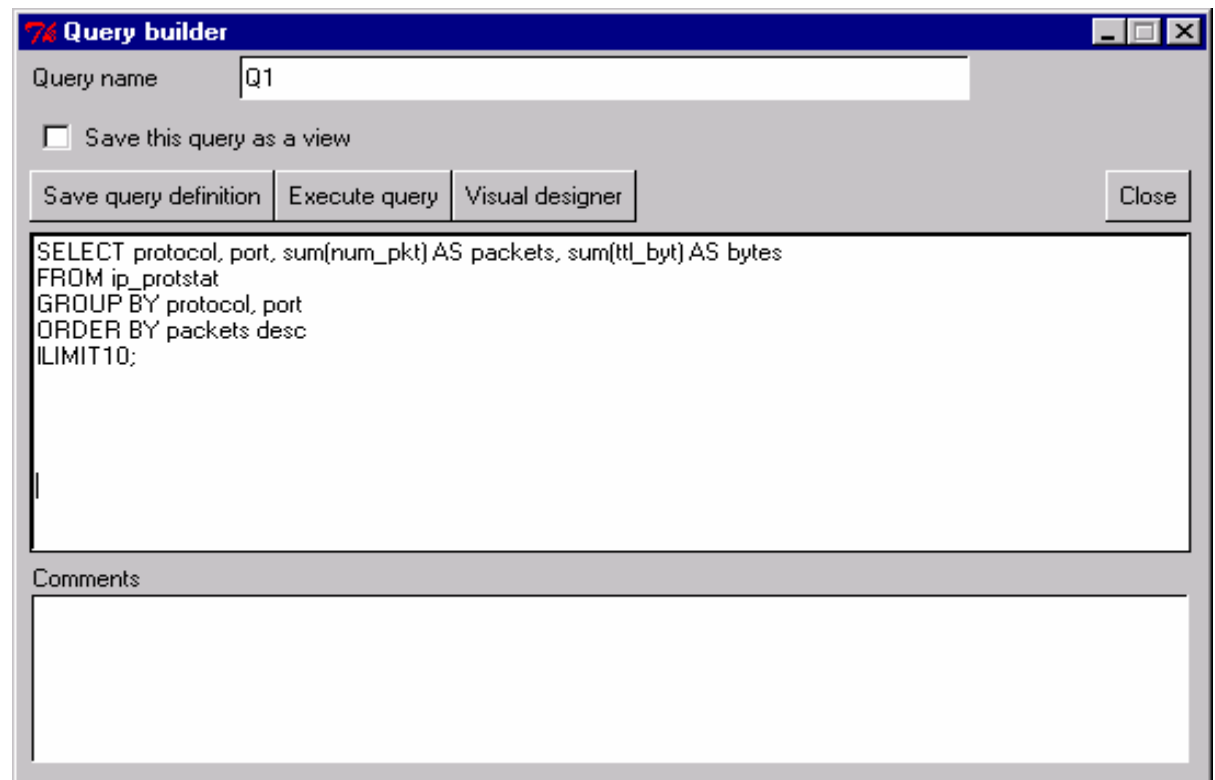
Database Environment



Sample Data Base Query

- **Query:** Report the top-10 protocol/port pairs ordered by number of packets transmitted

```
select protocol, port, sum(num_pkt) as packets, sum(ttl_byt) as bytes
from ip_protstat
group by protocol, port
order by packets desc
limit 10;
```



The screenshot shows a 'Query builder' window with a title bar containing a red '74' icon and the text 'Query builder'. The window has a 'Query name' field with 'Q1' entered. Below this is a checkbox labeled 'Save this query as a view' which is unchecked. There are three buttons: 'Save query definition', 'Execute query', and 'Visual designer'. A 'Close' button is in the top right corner. The main area contains a SQL query: 'SELECT protocol, port, sum(num_pkt) AS packets, sum(ttl_byt) AS bytes FROM ip_protstat GROUP BY protocol, port ORDER BY packets desc LIMIT 10;'. At the bottom is a 'Comments' section with a large empty text area.

74 Query builder

Query name: Q1

☐ Save this query as a view

Save query definition | Execute query | Visual designer | Close

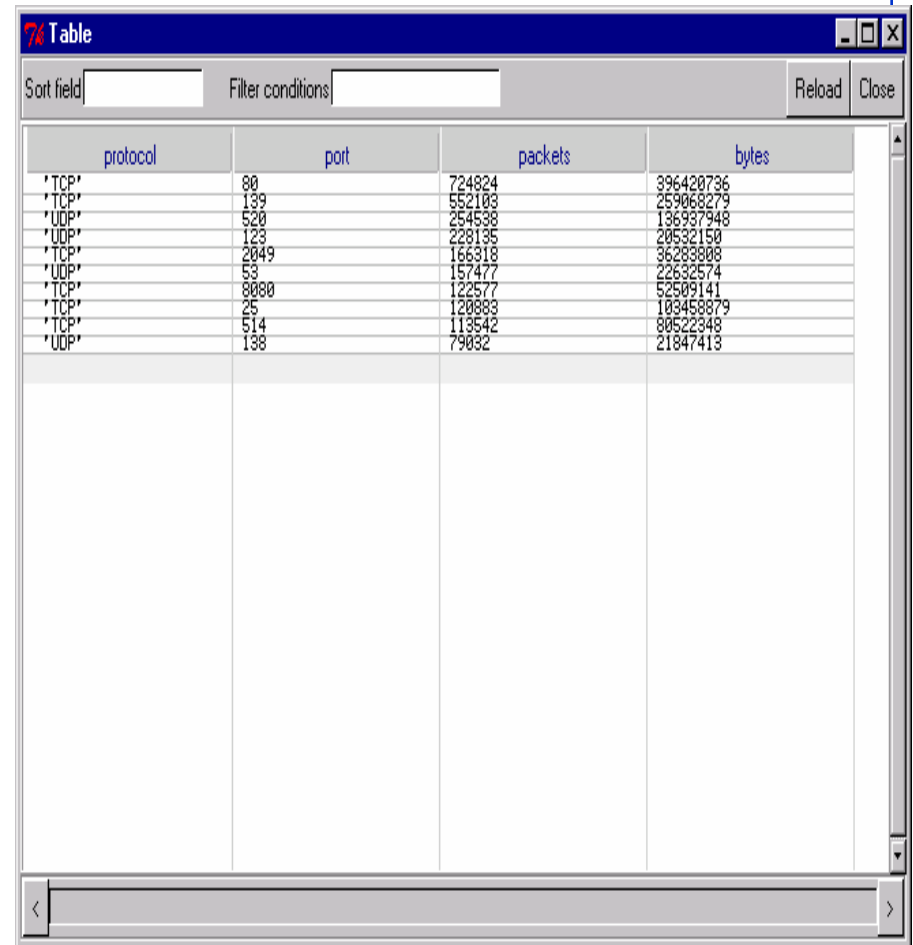
```
SELECT protocol, port, sum(num_pkt) AS packets, sum(ttl_byt) AS bytes
FROM ip_protstat
GROUP BY protocol, port
ORDER BY packets desc
LIMIT 10;
```

Comments

Response to Query

protocol	port	packets	bytes
'TCP'	80	724824	396420736
'TCP'	139	552103	259068279
'UDP'	520	254538	136937948
'UDP'	123	228135	20532150
'TCP'	2049	166318	36283808
'UDP'	53	157477	22632574
'TCP'	8080	122577	52509141
'TCP'	25	120883	103458879
'TCP'	514	113542	80522348
'UDP'	138	79032	21847413

(10 rows)



protocol	port	packets	bytes
'TCP'	80	724824	396420736
'TCP'	139	552103	259068279
'UDP'	520	254538	136937948
'UDP'	123	228135	20532150
'TCP'	2049	166318	36283808
'UDP'	53	157477	22632574
'TCP'	8080	122577	52509141
'TCP'	25	120883	103458879
'TCP'	514	113542	80522348
'UDP'	138	79032	21847413

IPEX Website

- Created by SLAC
- Traffic traces and time series plots of statistics
- HTML Interface
<http://ipexdata.research.telcordia.com/ipex/ipex.html>
- Restricted access for security - if you want access, send e-mail to Allen McIntosh <mcintosh@research.telcordia.com>
- Plan to extend access to data for interested groups in NMS

Current IPEX Analysis Status

Html Interface (<http://ipexdata.research.telcordia.com/ipex/ipex.html>)

IPEX Traffic Analysis

Connie Logg
Network Analyst, Stanford Linear Accelerator Center
cal@slac.stanford.edu, 650-926-2879
April 2002

This page provides links to the IPEX Traffic Analysis web pages.

Currently Available Reports

- **Daily Reports** include links for the specified day to: various IPEX Traffic results. Please note that not all probes have data for all days, hence you may get "The page cannot be found" messages. Access to the data for a probe for a day can be obtained by selecting the "Extracted Data" option in the days selection box.

Show the daily reports for:

IPEXPROBE.RESEARCH.TELCORDIA.COM	Jan	1	2002	2 days	Go
				2 days	
				10 days	
				6 weeks	
				Summary	
				Extracted Data	

Documentation

- **Port Definitions**


There are many sources of files which purport to define TCP and UDP port definitions. They often provide conflicting information. The [TCP](#) and [UDP](#) lists detail the origin of the port definitions used in this analysis. These are by no means "definitive", but are the ones chosen for port processing by SLAC.

Available Graphs


- Traffic volume: bytes/sec, bits/sec, packets/sec for
 - 2 days, 10 days, 6 weeks (the time period is easily changed).
 - Total IP, Total Non-IP, TCP, UDP, Miscellaneous IP
- Packet size histogram

Total IP Traffic


10 Days Bits/Second



Total Non-IP Traffic



Packet Size Distributions



Packet Size Histogram

Note: 50% of packets are < 100 bytes

Collector Security and Performance

- Collector Security
 - Probes and Database in DMZ
 - Protection from unauthorized access
 - Doing anything traverses two firewalls, an administrative headache
 - SSH access only
 - SSH error messages too cryptic (to deter bad guys?)
 - Deters good guys too
 - Web server still firewalled
 - Co-located with administrative function for another few weeks
 - Needs security check before general access can be enabled
- Collector Performance
 - Netra only good for 15 Mbits/sec

Phase I Future

- More collection sites
- Make other datasets available to NMS community

IPEX Phase II: Operational Measurement Infrastructure and Traffic Analysis

- **IPEX Goals:**

- insight into network performance through data studies and experiments
- Validation of results derived by analysis

- **New Focus: NMS community partnerships**

- Commercial sites not a viable test-bed for traffic measurement (other priorities in current business climate)
- Instead of concentrating on commercial sites, revise IPEX as ***NMS test-bed*** for traffic measurement (active and passive) and experiments, for relevant projects (Models, QoS,...)
- Results will be DOD focused and find broader commercial application

Ideas for IPEX Studies and Experiments

1. Sampling Issues in Measurements

- Sampling necessary because collection of complete traffic traces cannot keep up with increasing data rates
- What are good (perhaps, application-specific) sampling strategies?
- Implication for NMS
 - sampling strategies can make the difference between success and failure in achieving objectives (QoS monitoring, traffic characterization, performance control) for high-speed traffic
- Possible Studies
 - *Measuring QoS*

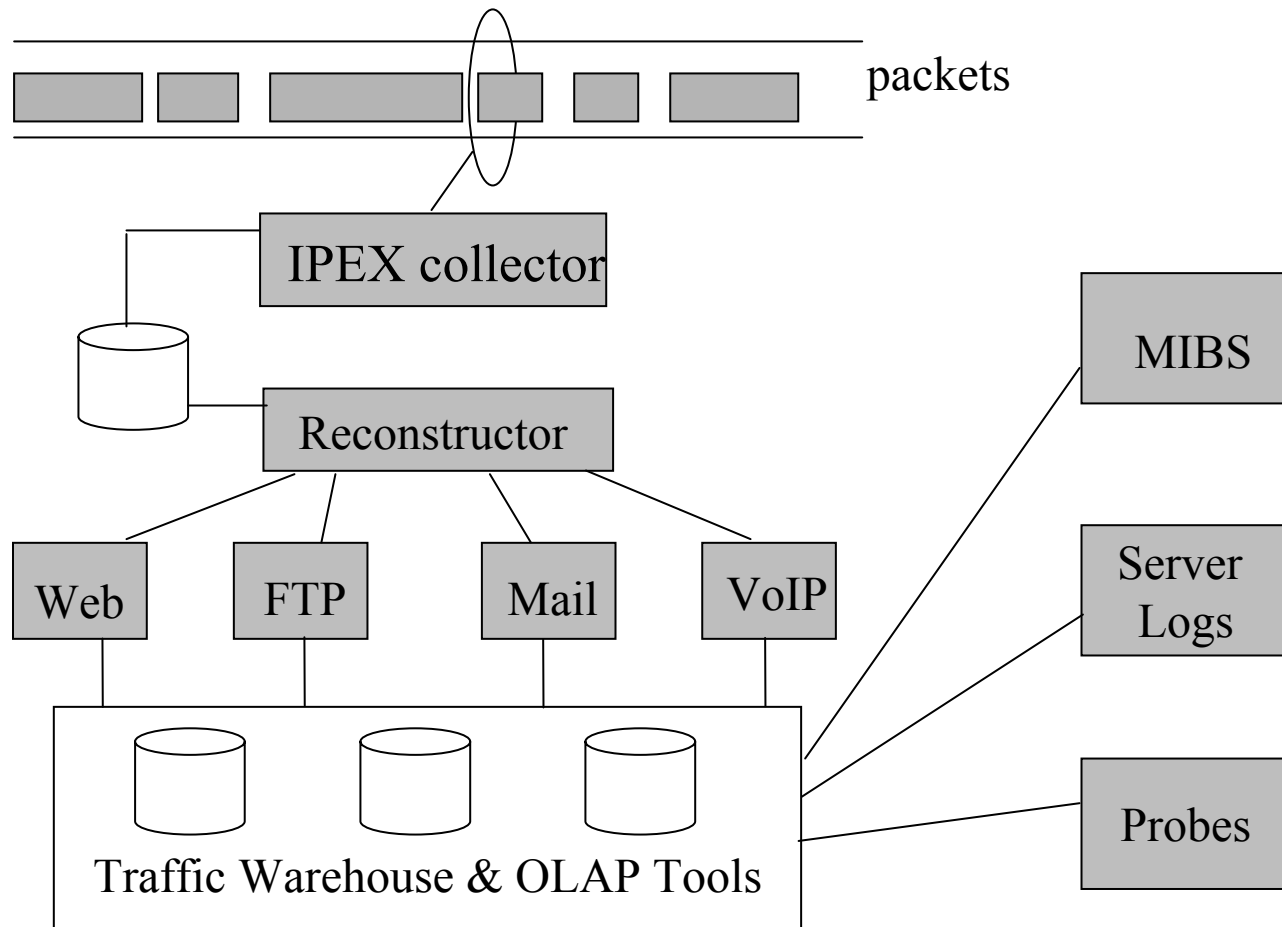
Determine sampling rate for acceptable accuracy of delay measurement, by examining variance as function of sampling rate (taking account of long-range correlation)
 - *Traffic models*

Comparison of models from detailed and 'sampled' versions of traffic trace, to determine measurement-resolution needed for deriving robust models

2. Data Mining for Analysis and Diagnosis

- Overall Aim
 - Exploit proven database techniques for uncovering patterns in data across multiple protocol layers, for traffic characterization and performance analysis
- Why needed?
 - Without correlation of application-level and packet-level information, information is often incomplete for
 - End-to-end service performance monitoring
 - Diagnosis of “root cause” of anomalies
- Complete data trace in IPEX would provide a *single* repository for analysis of *multiple* applications extending across multi-layer protocol stack
 - E.g., HTTP bandwidth and response times, Voice-over-IP jitter

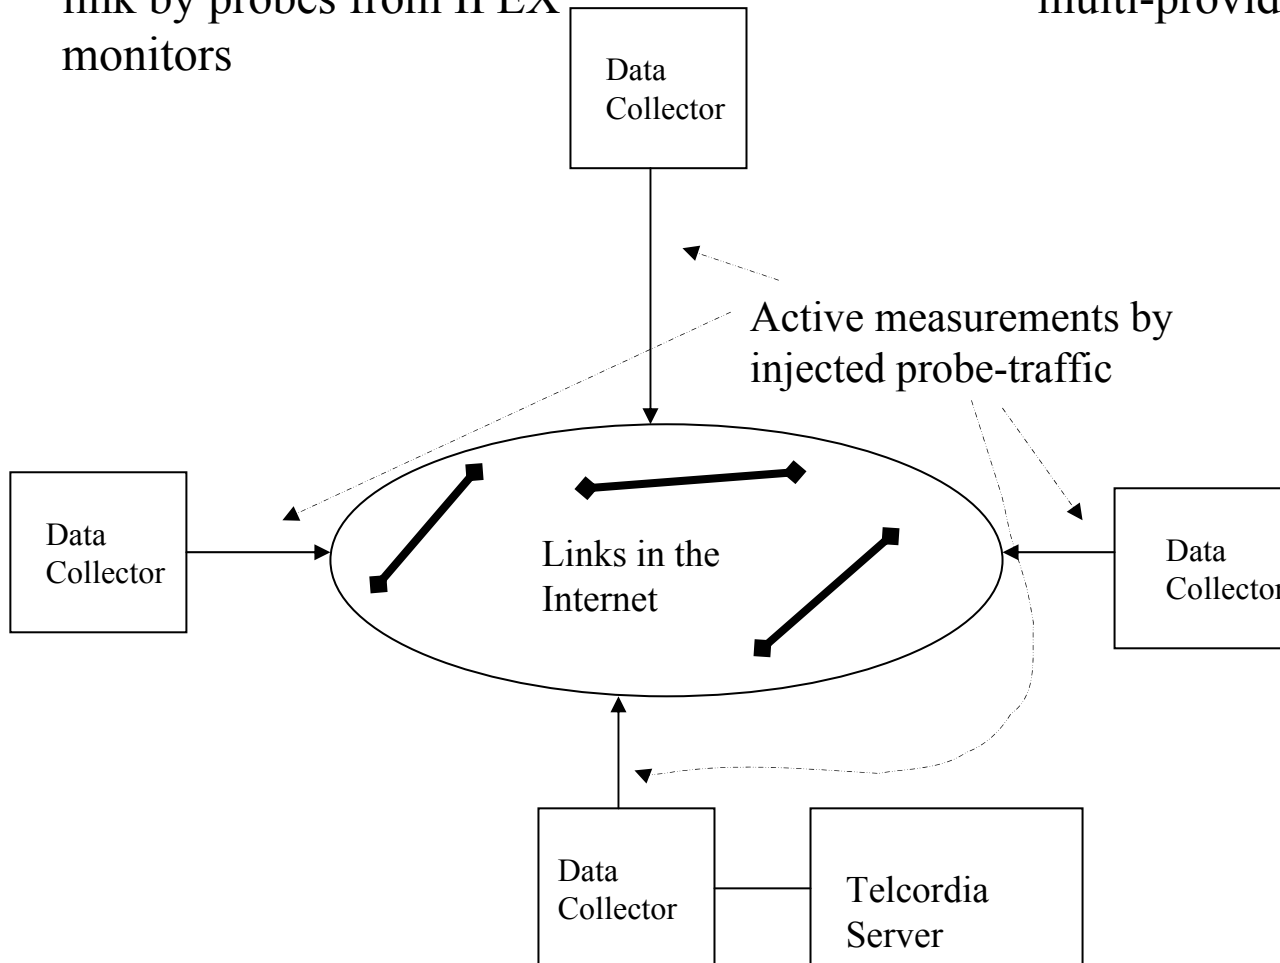
2. Data Mining for Analysis and Diagnosis (Cont'd)



3. Remote Monitoring of Internet Links by Active Probes

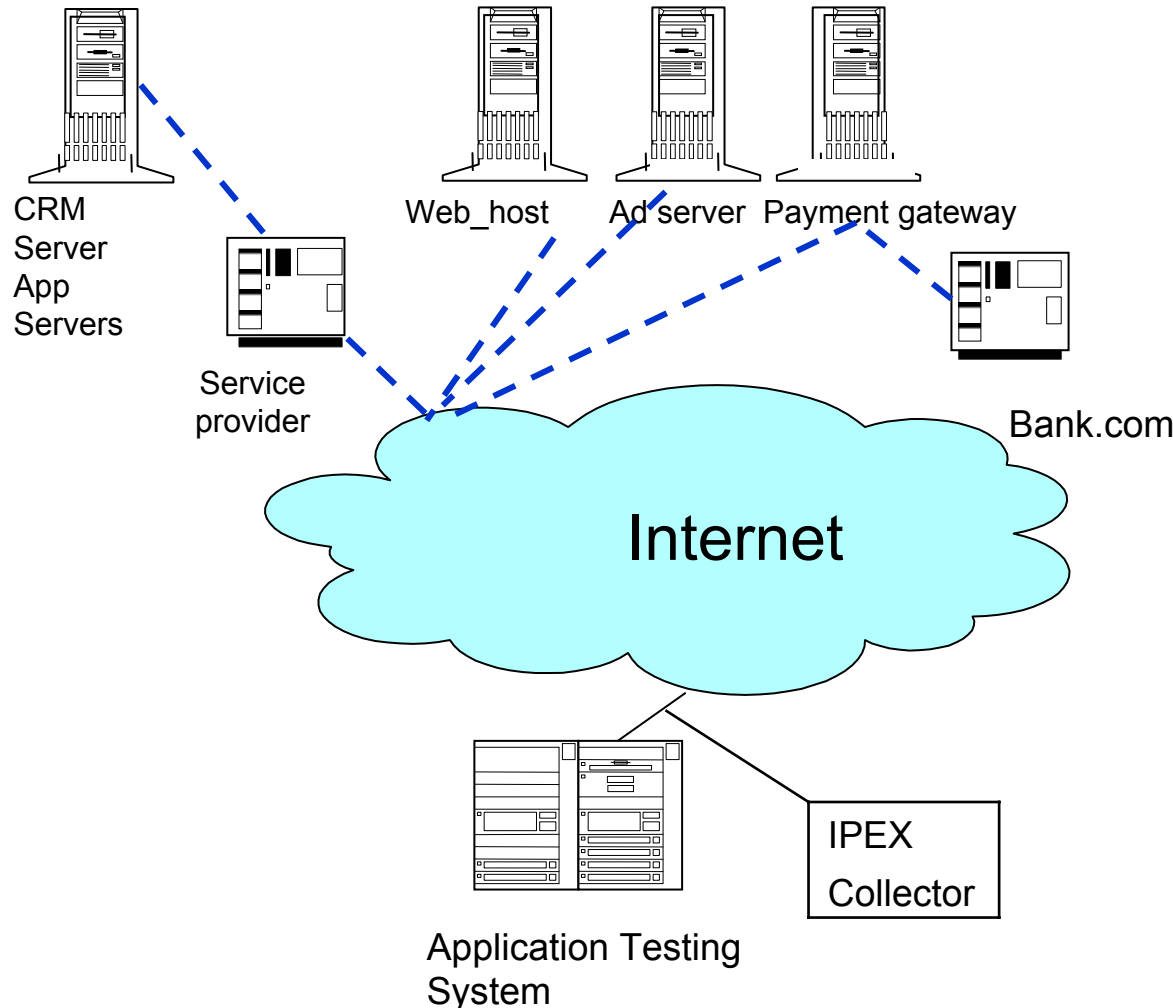
Surveillance of capacity and utilization of *remote* link by probes from IPEX monitors

“Intelligence” regarding remote bottleneck in multi-provider domain



4. Performance Testing of Web Service by Active Measurements

Cross-domain application testing by *synthetic* end-to-end transactions



Goal:

Distinguish between application-level and transport-level problems by analyzing packet traces and transactions data

Summary

- **Designed data collection infrastructure and database**
- **Started deployment**
- **Made data available to NMS community on the Web**
- **Next: Use infrastructure for experimentation**